

On Discounting Non-Marginal Policy Decisions and Cost-Benefit Analysis of Climate-Change Policy

Simon Dietz

- Department of Geography and Environment and Centre for Environmental Policy and Governance, London School of Economics and Political Science (LSE), Houghton Street, London WC2A 2AE, UK
- Tyndall Centre for Climate Change Research, University of East Anglia (UEA), Norwich, UK

Contact details: e-mail: s.dietz@lse.ac.uk
tel.: +44 (0) 207 955 7589
fax: +44 (0) 207 955 7412

With¹:

Alex Bowen

Bank of England, Threadneedle Street, London EC2R 8AH, UK

Cameron Hepburn

St Hugh's College, University of Oxford, St Margaret's Road, Oxford OX2 6LE, UK

Chris Hope

Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, UK

Nicola Patmore

Department for Environment, Food and Rural Affairs (DEFRA), Ashdown House, 123 Victoria Street, London SW1E 6DE, UK

Nicholas Stern

HM Treasury, 1 Horse Guards Road, London SW1A 2AE, UK

¹ All views expressed in this paper, and errors made, are Simon Dietz's.

VISIT...

LANZAROTE
Caliente.COM

Abstract

Monetary valuation of climate-change impacts, and the cost-benefit analysis of climate-change policy into which it feeds, has long been controversial. Writers in ecological economics have done much to illuminate its difficulties. For the purposes of this paper, the key difficulties of the cost-benefit approach are, first, discounting future costs and benefits at comparatively high rates, second, the often implicit assumption of perfect substitutability between natural capital and other forms of capital and, third, the treatment of uncertainty, much of which is very difficult to quantify in this case. On the other hand, alternative approaches to climate-change policy based on safe minimum standards are not without their weaknesses either. They are largely arbitrary and, as a result, arguably more politically unstable. Thus the in-principle need for cost-benefit analysis will not diminish.

In this paper, we consider how to develop the theory and practice of cost-benefit analysis as applied to climate-change policy, with a keen eye on its main weaknesses. This in particular demands that we look closely at the theory of inter-temporal aggregation: of discounting. We present an expected-utility approach to cost-benefit analysis. Our central premise is simple: given that unabated greenhouse gas emissions could potentially entail large, irreversible costs on the global economy and on welfare (formally, if the impacts of climate change are non-marginal), the simplified approaches to cost-benefit analysis and discounting practised in project appraisals the world over are inadequate. Instead, an expected-utility approach is required that calculates utility integrals along a range of possible paths. Expected utility must be calculated separately along business-as-usual paths compared to stabilisation paths, or else significant errors could be induced in the estimation of costs and benefits. This in turn requires that we find alternative ways of expressing the present value of the total and marginal costs of greenhouse gas emissions. We use the PAGE2002 integrated assessment model to develop some of our theoretical contribution empirically.

1. Cost-benefit analysis of climate-change policy

Monetary valuation of climate-change impacts, and the cost-benefit analysis (hereafter CBA) of climate-change policy into which it feeds, has long been controversial. Since the first wave of studies was published in the early-mid 1990s (e.g. Cline, 1992; Fankhauser, 1995; Maddison, 1995; Manne *et al.*, 1995; Nordhaus, 1991, 1994; Peck and Teisberg, 1992; Plambeck and Hope, 1996; Tol, 1997), the same conclusion has, with occasional exceptions (notably Cline, 1992), been drawn: the monetary impacts of future climate change, discounted back to the present day, are not as large as many would have us believe and, when pitted against the cost of reducing greenhouse gas (GHG) emissions, the optimal policy will bring about but a small reduction in emissions and global mean temperature. Stabilising atmospheric concentrations of GHGs, the ‘holy grail’ of the scientific community, is typically sub-optimal, at least within the next century or so.

Of course, these conclusions, and the cost-benefit approach more generally, have been subject to criticism, with the ecological economics community playing an important role. Of the many arguments made, three are of particular relevance to this paper:

1. The ethics behind, and results of, discounting future costs and benefits at comparatively high rates (e.g. Azar and Sterner, 1996; Broome, 1992; Cline, 1992);
2. The implicit assumption of perfect substitutability between natural capital and other forms of capital made in the most aggregated cost-benefit models (highlighted by Neumayer, 1999²);
3. The treatment of uncertainty (e.g. Azar, 1998; Neumayer, 1999 and various contributions to *Climatic Change*, 56(3)).

Another important set of criticisms has been levelled at the welfarist ethical framework underpinning CBA. Along this line of reasoning, it may be unethical to impose environmental damage on future generations, as compensation by way of increased consumption of non-environmental goods, and an increased stock of other forms of capital, is inadequate (e.g. Barry, 1991; Spash, 1994). Legitimate as these concerns are (see Stern, 2006, chapter 2, for a discussion and implications for climate-change policy appraisal), in this paper I adopt an avowedly welfarist perspective.

The ‘tyranny’ of discounting³

² And developed in terms of substitution in the utility function by Gerlagh and van der Zwaan (2002).

³ To borrow the term from Pearce *et al.* (2003).

First, because the climate-change externality occurs with a lag but thereafter persists, the benefits of avoided climate change accrue long after many of the costs must be paid. Discounting at comparatively high rates will reduce the future benefits of avoided climate change to the extent that even fairly significant reductions in income and welfare in the far-off future will be worth absorbing in the process of development. Cline (1992) provides one of the earliest and best discussions: in it, he for instance demonstrates the simple calculation that \$1, received in 200 years' time, is worth one half of one millionth of one cent today when discounted exponentially at 10% per annum.

Capital substitutability

Second, the cost-benefit approach aggregates impacts on a wide range of environmental functions and goes on to aggregate the monetary value of these with the total value of consumption and investment in the economy in an additively separable framework. As such it makes an implicit assumption of perfect substitutability between natural capital and other forms of capital. Optimal climate-change policy does not involve strong emission abatement, because future generations can be compensated for losses in natural capital assets through investment in other forms of capital and increased consumption of flows from the latter. Starting with the work of Daly and Costanza (Daly, 1992, 1996; Daly and Costanza, 1992), ecological economists are familiar with the counter-notion that at least some portion of the world's natural capital stock is indispensable and thus the relationship between natural capital and other forms of capital would in many cases better be described as complementary.

The difficulty for the proponents of such an argument lies in determining precisely where the threshold lies, below which we cannot continue to deplete the global climate resource through emissions in excess of natural sink capacity. It seems quite clear that, at a basic level, the global environmental and ecological system, which provides us with life support functions such as stable and tolerable climatic conditions, cannot be substituted. What volume of GHGs can be emitted before any threshold is reached that would put these functions at risk is a much more difficult question.

Eric Neumayer (1999) brings the comparative importance of the substitutability issue, versus the discounting issue, into sharp focus. For him, critics of 'excessive' discounting of the future are missing the point, which is that any discounting of future consumption streams, premised on subtracting the monetary value of climate-change impacts from projected growth in aggregate consumption, is implicitly based on the assumption of perfect substitutability.

Third, the question arises as to how much confidence we can place in monetary estimates of the costs and, especially, the benefits of GHG emission reductions when the uncertainties are so large? That large uncertainties are associated with every step in the chain that links long-run trends in population, economic growth and technology, through the entire process of anthropogenic climate change, to impacts on the natural environment, economies and societies is undisputed. As Tol (1999) notes, much of this uncertainty is irresolvable, because it relates to inherently stochastic processes and because we are in many cases attempting to predict trends into the far-off future.

Yet the models on which CBA of climate-change policies is carried out – so-called Integrated Assessment Models (IAMs) – have not always chosen to tackle this uncertainty directly. Deterministic approaches to modelling largely finesse the uncertainties, focussing often on ‘best guesses’ for their many parameters. For the direct purpose of valuing the social costs of GHG emissions and using CBA to determine efficient policies (in a one-time decision context), what is required is a stochastic modelling approach, allied to measures of cost and decision rules that apply the axioms of the economics of risk, beginning of course with von Neumann and Morgenstern (1944).⁴ Doing just so is one of the primary foci of this paper. There are some examples of the application of a probabilistic approach to directly estimate climate-change impacts in the literature, such as Mastrandrea and Schneider (2004), Plambeck *et al.* (1997), and Tol (1999, 2003). However, they have generally stopped short of a full expected-utility analysis, or else they have focussed on very specific aspects of the issue. In this paper, we aim to clearly state an expected-utility approach and to illustrate the consequences of adopting it, relative to more simple CBA techniques.

But in any case climate change would not best be described as a situation of ‘risk’. Many elements of the overall uncertainty attached to the consequences of GHG emissions would better be described as situations of uncertainty, in the Knightian (Knight, 1921) or ‘hard’ sense. Whether CBA remains a meaningful decision-support tool then depends on our ability to define subjective probability distributions over the relevant variables. This is indeed the necessary core of any stochastic modelling approach. If our ability to do so is very limited, then so is the applicability of CBA (e.g. Neumayer, 1999).

⁴ Of course there are a number of dimensions to the uncertainty problem, including learning. Because some of the uncertainty surrounding future climate-change impacts is due to a poor understanding of the processes in question, and if, as seems reasonable, we expect to obtain better information about these processes in the future, then the measurement of quasi-option values suggests itself. While such dimensions are undoubtedly of great importance, in this paper we focus on the more direct task of valuing climate-change impacts, with a view to undertaking CBA.

Is CBA irrelevant?

In light of these concerns, the natural approach for ecological economics to take towards climate-change policy appraisal is an application of safe minimum standards, known in the climate-change context as the ‘safe corridor/landing’ approach or the ‘tolerable windows’ approach. Allocative efficiency is of secondary concern, though the cost of GHG emission reductions does enter the calculus in an *ad hoc* manner. On the other hand, safe minimum standards are arbitrary and, largely as a result, politically unstable too. Admittedly, CBA is far from insulated from political pressures, given that key decisions on aggregation across individuals in different regions and time-periods are fundamentally rooted in ethics (see below). Yet within the ‘corset’ of CBA, such ethical judgements ought to be made systematically and clearly, so, if we at least accept that the ethical underpinnings of welfarism will often command broad appeal (Hubin, 1994), there will always be an in-principle requirement for innovation in CBA.

Hence in this paper I in effect ask whether we can push social CBA forward, so that the concerns highlighted quite effectively by ecological economics can be taken on board within a welfare-economic framework? Doing so, however, demands a departure from the conventional simplifications of discounting applied in project appraisals the world over. Section two develops the theory here, and derives alternative measures of the marginal and total cost of GHG emissions. Section three introduces the empirical model, the PAGE2002 IAM (Hope, 2006), used in this paper to demonstrate our theoretical insights. PAGE2002 has two features that make it especially suited to the present research problem. First, it is a stochastic model, producing a probability distribution of climate-change impacts based on a Monte-Carlo procedure. Second, in the spirit of Nordhaus and Boyer’s (2000) pioneering estimates, it includes the possibility of catastrophic climate impacts (or ‘large-scale discontinuities’, in the language of the IPCC in 2001: McCarthy *et al.*, 2001). This is crucial, since I would argue that modelling such processes arguably simulates the consequences of depleting irreplaceable or ‘critical’ natural capital. More specifically, examples include a significant slow-down in the thermohaline ocean circulation, large-scale melting of the Greenland and West Antarctic ice sheets and so on (McCarthy *et al.*, 2001). Though highly uncertain, the natural, economic and social cost of such discontinuities is likely to be very high indeed and this is reflected in the way Nordhaus and Boyer (2000) and Hope (2006) model it. Strong convexities in the damage functions linking increases in global mean temperature with impacts on economic sectors, ecosystems and human mortality might also be interpreted as the consequence of the depletion of critical natural capital, in the former case critical as a

factor of production and in the latter case critical as an argument in the utility function of a representative consumer.⁵

Thus the question as to whether CBA can be applied where some of the natural capital stock is non-substitutable may not be answerable with a simple ‘yes’ or ‘no’. Instead, I would argue that it depends on our ability to assign subjective probabilities to the consequences of reaching and exceeding critical thresholds. Ultimately, the consequences of depleting critical natural capital must be felt through losses in income and welfare. So discounting must still be at issue here, because as Cline’s (1992) simple example above shows, losses in global income in the far-off future – even precipitous losses – will be worth little in present value terms with a high discount rate. For instance, taking projections of growth in gross world product to 2200 from one scenario (the IPCC A2 scenario, extrapolated to 2200 by Hope, 2006), a one-time 50% loss in GWP in 2200, amounting to an undiscounted real value of \$1.06 quadrillion (year 2000 prices), is worth just \$5.6 million in present value terms, discounting exponentially at 10% per annum.

Section four presents our empirical results and section five provides a discussion. Although the results and conclusions of this paper are relevant for CBA in general, we focus most often on the benefits of GHG emission reductions: on monetising the benefits of avoiding climate change.

2. Discounting non-marginal policy choices

Towards endogenous discounting

Why will simplified, project-appraisal CBA not suffice in the case of climate-change policy? The basic answer is that the potential changes in income and welfare that major, global choices over whether or not to stabilise emissions might entail are too large to be handled using the tool of discount rates, a method for analysing changes that are marginal with respect to the whole growth path. We can understand this point with recourse to the concept of the social discount factor, which describes the social price of consumption in the future compared to consumption today. Presenting the analysis in terms of social discount factors is very constructive in this case, because it opens the analysis to the insight that divergent paths of future per-capita consumption lead, all else equal, to divergent sets of social discount factors, and, in turn, to divergent average social discount rates.

⁵ In the case of complementarity in the utility function, a more direct approach would hence be to actually specify environmental amenities as a separate argument, as in Tol (1994).

The social discount factor, $\lambda(t)$, is the marginal social value of a unit of consumption (or output) in year t relative to now: it is a concept defined with respect to a path of consumption per capita $[c(t)]$. Thus the relative worth of an extra unit of consumption in time t , to a unit now, is likely to depend on $c(t)$ and $c(0)$ (as well as other possible levels of c at different times). If $c(t)$ is bigger than $c(0)$ then we might expect λ to be less than one, given diminishing marginal utility of consumption (postponing for the moment a discussion of pure time discounting). And if c were rising around (t) then λ would be falling. Thus the social discount rate, the rate of change of λ , would in this case be positive. Formally, the social discount factor is linked to the social discount rate by the following identity:

$$\lambda(t) = \frac{1}{(1 + s(t))^t} \quad (1)$$

Where s is the social discount rate. The point here is that the social discount rate is likely to be very sensitive to differences across paths where we are considering the possibility of severe climate-change impacts on the economy. Indeed climate-change impacts could be so severe in some parts of the world that c is falling around t and so λ would be rising and s would be falling. Thus simple CBA with an exogenous discount rate is not an appropriate tool.

In recent years, economists have developed analyses that account for exogenous uncertainty in future consumption. These have given rise to the idea that uncertainty over future consumption can be tackled via the application of an exogenous, certainty-equivalent, time-declining discount rate (e.g. Gollier, 2001, 2002a, b; Groom *et al.*, 2005; Newell and Pizer, 2003; Weitzman, 1998, 2004). The difficulty is that this approach depends on the assumption that uncertainty over future consumption is independent of the policy choice being contemplated. This may not be the full story here, because the policy decision in question may itself lead to large differences in consumption paths, or more pertinently large differences in the probability distribution of future consumption paths. Hence we are faced with endogenous uncertainty, as well as exogenous uncertainty.⁶

Instead we have to go back to the underlying objective against which consumption paths are evaluated over time (and from which discount rates are derived); generally an integral of utility of consumption. This is not, of course, to abandon a basic cost-benefit calculus but simply to avoid using an often helpful short-cut, the discount rate, which is not valid in this context of non-marginal changes.

⁶ Exogenous uncertainty undoubtedly does still exist, but in a stochastic modelling framework we finesse this point by assuming that all uncertainties have been parameterised.

To further develop this approach, we first need to specify our social welfare function. We begin by specifying the utility of the representative consumer as a simple, iso-elastic function of consumption per capita in the standard way:

$$u(t) = \frac{c(t)^{1-\eta}}{1-\eta} \quad (2)$$

Where u is utility per capita, c is consumption per capita and η is (the negative of) the elasticity of marginal utility with respect to consumption. Discounted per-capita utility is given in continuous time by:

$$w = \int_{t=0}^{\infty} u(\tau) \exp^{-\delta\tau} d\tau \quad (3)$$

Where w is discounted utility per capita and δ is the pure rate of time preference or utility discount rate. Finally, with exogenous population, we would automatically weight by population to derive our aggregate social welfare or objective function:

$$W = \int_0^{\infty} N(\tau) u(\tau) \exp^{-\delta\tau} d\tau \quad (4)$$

Where W is social welfare and N is population. Under non-marginal uncertainty, every possible outcome i of a climate-change policy choice, which has associated with it a probability p , should be evaluated with the social welfare function in equation (4). If we were also able to measure consumption and the impacts of climate change in different world-regions, then W should be calculated separately for each region. This is an item for further work. We are now able to compute expected utility for the probability distribution of consumption paths simulated to follow from a policy choice:

$$E(W) = \sum_{i=0}^N p_i W_i \quad (5)$$

Where E denotes the expectation, for possible outcomes $i=1, \dots, N$. Thus we are discounting endogenously, and every possible consumption path flowing from the policy choice will have a unique set of discount factors, and hence a unique average discount rate, associated with it. Overall, the deterministic discount factor at time t is replaced by its expectation, the (endogenous) certainty-

equivalent discount factor (the rate of change of which is of course the certainty-equivalent discount rate):

$$\lambda(t) = \sum_{i=1}^N p_i c(t)^{-\eta} \exp^{-\tilde{\alpha}} \quad (6)$$

Deriving the total and marginal costs of GHG emissions in an expected-utility framework

Performing the analysis in this way renders the task of expressing the total and marginal costs of GHG emissions more difficult. Under project-appraisal CBA, we would simply calculate the social costs, in each period, of GHG emissions and determine their present value by discounting at an exogenous rate. In undertaking full CBA, we would of course include the costs of emission reductions, discounting the net benefits of these reductions. In this case, however, we are left with the task of expressing expected discounted utility, in utils, in a way that is useful to policy.

To calibrate total costs, we apply the concept of the ‘balanced growth equivalent’ (BGE), developed by Mirrlees and Stern (1972). The BGE essentially measures the utility generated by an actual, simulated consumption path in terms of the consumption today that, if it grew at a constant rate, would generate the same utility (yield an equal utility integral). Formally, the change in the BGE is a natural commodity measure of welfare that, in general terms, expresses changes in future consumption due to policy in terms of the percentage change in consumption (along a steady-state growth path), now and forever, that is equal to the changes that are forecast to follow from that policy. In a one-sector growth model with natural growth α and consumption C at time t , we want to calibrate welfare from the path $[C(t)]$. If this is equivalent, in welfare terms, to the balanced growth path yielding consumption $\gamma \exp^{\alpha t}$, then γ is the BGE of $[C(t)]$.

Adopting BGE differences as our measure of total cost requires that we find a means to extrapolate simulated consumption paths during the modelling horizon (in PAGE2002 this is a 200-year period from 2001 to 2200) into perpetuity. We approximate this by assuming that every one of the model’s simulated consumption paths grows at a constant, arbitrary rate beyond the modelling horizon:

$$W = \sum_{t=0}^T N(t)u(t) \exp^{-\tilde{\alpha}} dt + \left(\frac{N(T)u(T)}{\delta - g(1-\eta)} \right) \exp^{-\tilde{\alpha}} \quad (7)$$

Where g is the arbitrary rate of growth of per-capita consumption beyond year T , 2200. The practical implication of this is that we assume no additional impacts of climate change after 2200. This is clearly unrealistic and, as we move the emission of GHGs further and further into this century or even the next, the modelling horizon becomes less and less adequate to deal with their full social cost. On the other hand, we assume the impacts that give rise to each consumption path with probability p drive a permanent shortfall in consumption relative to a world without climate change. That is, there is hysteresis of climate impacts into perpetuity, the significance of which to the overall evaluation depends on the rate of fall of the social discount factors on the path. The faster they fall, the less it matters.

We must now find the BGE of a given consumption path. We do so iteratively by specifying an initial level of per-capita consumption γ that, growing at an arbitrary constant annual rate, yields the same amount of utility as the path:

$$u(\gamma) = \frac{c(\gamma)^{1-\eta}}{1-\eta} \quad (8)$$

$$\text{and} \quad W = \sum_{t=0}^T N(t)u(\gamma + gt) \exp^{-\alpha} dt + \left(\frac{N(T)u(\gamma + gT)}{\delta - g(1-\eta)} \right) \exp^{-\alpha} \quad (9)$$

Where g is again the arbitrary rate of growth of per-capita consumption, set in this case equal to g in equation (7). Note that the arbitrary choice of g does not affect the relative differences between BGEs and therefore the calibration of the total cost of GHG emissions. Instead it simply scales the initial amount of consumption per capita required to approximate the actual consumption path. As g decreases, the BGE per-capita consumption level γ must increase and *vice versa*.

In terms of the marginal cost of GHG emissions, the thought experiment that we must undertake is as follows: what fall in initial aggregate consumption equals the difference in expected utility between a path of GHG emissions and a second path of GHG emissions, with an additional pulse of one tonne of carbon in the initial year? Letting $\hat{\cdot}$ denote the former path and $\tilde{\cdot}$ denote the path with an additional one tonne of carbon, we can calculate the difference in expected utility:

$$\Delta E(W) = E(\hat{W}) - E(\tilde{W}) \quad (10)$$

Then we can calculate the initial level of utility per capita that would yield the same amount of expected discounted utility along the path with an extra one tonne of carbon:

$$\tilde{u}(0) = \frac{E(\hat{W}) + \Delta E(W)}{N(0)} \quad (11)$$

Finally, the marginal damage cost of carbon is the difference between initial aggregate consumption with an extra tonne of carbon and without:

$$\tilde{c}(0) = [(1 - \eta)\tilde{u}(0)]^{\frac{1}{1-\eta}} \quad (12)$$

$$\text{So } MDC(0) = N(0)\hat{c}(0) - N(0)\tilde{c}(0) \quad (13)$$

Where MDC is the marginal damage cost of carbon.

How much to discount the future?

Having set out our framework, it remains to comment on the choice of two central parameters for evaluation: (1) the elasticity of marginal utility of consumption and (2) the pure rate of time preference or utility discount rate. It is beyond the scope of this paper to provide a thoroughgoing discussion (the interested reader might begin with recent reviews by Hepburn, 2006, and Stern, 2006). Nevertheless, the results we present show that these two choices are quite critical for the overall results (as has been noted, in the wake of the Stern Review, by e.g. Nordhaus, 2006, and Yohe, 2006).

Turning first to the elasticity of marginal utility of consumption, the typical approach is to calibrate this parameter based on empirical estimates, derived from a variety of sources such as individual saving behaviour and social inequality aversion, the latter as revealed for example by the structure of a tax system. Most estimates point to some degree of concavity in the utility function. Although a survey by Stern (1977) raised the possibility of a wide range of elasticities running from 0 to 10, with elasticities around 2 most supported (albeit with a strong warning against placing too much confidence in the range for the purposes of CBA), more recent work by Cowell and Gardiner (1999) concludes 0.5 to 4 is reasonable. Pearce and Ulph (1999) would put η in the range of 0.7 to 1.5 and OXERA (2002) note that even an elasticity as low in the range as 2 implies very high – in their view implausible – inequality aversion.⁷ This line of argument tends to lead to a convenient assumption that η is unity, so that the general form iso-elastic utility function in equation (2) simplifies to the natural logarithm of

⁷ In a simple framework where a ‘rich’ individual enjoys ten times as much consumption as a ‘poor’ individual, if η is 2, then an additional unit of consumption to the poor individual is worth 100-times more than it would were it to accrue to the rich individual.

consumption per capita. In this case, the contribution to the social discount rate of increases in per-capita consumption is simply equal to the rate of growth in per-capita consumption. Note that the general approach to specifying a single, iso-elastic utility function is rather restrictive, in the sense that it assumes the structure of preferences to avoid risk is the same as preferences to smooth consumption. Further research will invoke alternative preference specifications to relax this assumption (e.g. Epstein and Zin, 1989; Kreps and Porteus, 1978, Selden, 1979).

The second choice, over the pure rate of time preference or utility discount rate, is much contested. In the briefest of treatments, we can distinguish between two basic approaches (e.g. Arrow, 1995). One may be characterised as ‘descriptive’ and seeks to calibrate δ based on the observed time preference or impatience of individuals (see Pearce and Ulph, 1999), either implicitly or explicitly taking into account life chances and their effect on the valuation of future consumption. In some cases, δ is solved as the difference between observed rates of saving and/or interest and the product of consumption growth and its marginal utility (e.g. Manne *et al.*, 1995; Nordhaus, 1994). The other may be characterised as ‘prescriptive’ and argues that δ should be prescribed based on an ethical evaluation of the relative weight to be apportioned to the consumption of future generations, irrespective of their wealth. On this basis, it becomes very difficult to resist the argument that future generations merit equal ethical weight (linked to the utilitarian tradition by Broome, 1992), so that δ is calibrated on the (probably tiny) chance that future generations will not exist to enjoy consumption, due to some sort of global catastrophe⁸ (see Stern, 2006, for a longer discussion of this position, which, as is often noted, has a rich heritage in economics beginning with Ramsey, 1928). The most persuasive counter-argument to this is that such low time preference implies that the current generation is required to save excessive amounts in support of the far-off future (e.g. Arrow, 1995), not in itself an apparently ‘ethical’ outcome. This point is perhaps too coarse, however, since it can be circumvented by increasing the concavity of the utility function η (Asheim and Buchholz, 2003), requires a more careful focus on other variables relevant to saving, such as the lifetime of capital equipment, and would have to apply to all public investments today, even though in practice few provide a direct stream of consumption over such a long time-scale as GHG emission reductions (i.e. greater time preference might reasonably be permitted for other public investments that pay back over shorter time-periods, while still being consistent with an impartial approach to very long-run investments).

⁸ More precisely, δ becomes the annual risk of wiping out humanity, as represented by a Poisson process $\exp^{-\delta T}$ over time T .

3. Our model

To demonstrate our theoretical arguments empirically, we employ the PAGE2002 IAM. We have already commented on the two main advantages of PAGE2002 in section one: first, that it is a stochastic model, designed to generate a probability distribution of consumption paths under climate change and climate-change policy; and, second, that it is flexible enough to include the small risk of catastrophic climate impacts and to simulate strong convexities in the damage function describing non-catastrophic impacts. The model is comprehensively described in Hope (2006) and, unless otherwise stated, we make no changes to the parameters reported therein. The primary features of the model are as follows:

- Emissions of carbon dioxide and methane (including natural emissions stimulated by climate change in the model), as well as SF₆, and other GHGs that contribute to background radiative forcing (e.g. NO₂ and (H)CFCs);
- The accumulation of anthropogenic GHG emissions in the atmosphere and resulting radiative forcing;
- Regional temperature increases arising from the difference between greenhouse warming and regional cooling brought about by sulphate aerosols;
- Nonlinearity and transience in the damage caused by increases in regional mean temperature, modelled as a polynomial function above a time-varying tolerable level of temperature change;
- Region-specific impacts of gradual climate change in two main sectors – (1) ‘economic’ impacts on ‘market’ sectors of the economy like agriculture and energy and (2) ‘non-economic’, direct welfare impacts on ecosystems and human health;
- Adaptation to climate change, brought about by investment, that increases the time-varying tolerable level of climate change before economic losses (category (1) above) arise and reduces the intensity of impacts on both economic and non-economic sectors above the tolerable level. Adaptation is also region-specific;
- The possibility of future catastrophic impacts. When global mean temperature rises to high levels (an average of 5°C above pre-industrial levels), the chance of large losses in regional GDP in the range of 5-20% begins to appear. This chance increases by an average of 10% per °C rise in global mean temperature beyond 5°C.

31 key inputs to the model are stochastic. A third key feature of the model is that it is essentially meta-analytical, insofar as the density functions for the various stochastic parameters are calibrated in their minima, modes and maxima to the underlying scientific and economic literatures on climate change. Thus PAGE2002 has in the past been shown to produce results, for instance with respect to the

marginal damage cost of carbon (conventionally calculated), close to the centre of a range of peer-reviewed studies (cf. Hope, 2006 with Tol, 2005).

For the purposes of this paper, we run the model with two different climate scenarios. These are described in more detail in Stern (2006, chapter 6). Our ‘baseline-climate’ scenario is designed to give warming consistent with the IPCC *Third Assessment Report* in 2001 (IPCC, 2001). Our ‘high-climate’ scenario adds to this a risk of amplifying natural feedbacks in the climate system. This is based on recent studies showing that there is a risk of additional feedbacks, such as weakening carbon sinks (Friedlingstein *et al.*, 2006) and natural methane releases from wetlands and thawing permafrost (e.g. Gedney *et al.*, 2004). The high scenario increases the probability of larger temperature changes. In 2100, the baseline scenario produces a mean warming of 3.9°C relative to pre-industrial and a 90% confidence interval of 2.4 – 5.8°C for the IPCC A2 emissions scenario (Nakicenovic and Swart, 2000, adapted by Hope, 2006) used in this exercise. In contrast, the high scenario gives a mean warming of 4.3°C in 2100 and a 90% confidence interval of 2.6-6.5°C. Unlike Stern (2006), however, we keep the structure of the analysis simple by confining ourselves to simulations with all three impact sectors (economic impacts, non-economic impacts and the risk of catastrophe).

We set PAGE2002 to produce 1000 runs of aggregate GWP, between 2001 and 2200, net of the monetary cost of climate-change damage and the cost of adaptation to climate change. Clearly, our analysis is based on a number of assumptions. Those particular to this paper are reported in table 1. A final remark is that, in order to ensure converge of the expected-utility integral on any given simulation draw where the elasticity of marginal utility of consumption is 1, we cap impacts at 99% of GWP (see section five for a discussion of this). In effect, we specify a crude s-curve in impacts. For now, with our analysis confined to the global level, this cap is highly unlikely to come into effect, except for one or two draws in the high-climate scenario. If the analysis is taken to the regional level, it may prove more important.

Table 1. Summary of assumptions.

<i>Parameter</i>	<i>Assumption</i>
Economic growth	Based on IPCC A2 scenario (see Nakicenovic and Swart, 2000), converted to PPP and extrapolated to 2200 in Hope (2006).
Population growth	As above
Saving rate	0.2, constant
Elasticity of marginal utility of consumption	1.0; 1.25; 1.5; 2.0
Pure rate of time preference	0.1; 0.5; 1.0; 1.5
Constant rate of growth of consumption per capita after 2200	1.3% p.a. (the average between 2001 and 2200 without climate change)

4. Results

Basic estimates of total costs

We begin with a brief summary of the undiscounted cost of climate change through time (figure 1a, b), starting in 2050. Figure 1a traces the mean percentage loss in GWP, as well as the range of losses from the 5th to the 95th percentile, for the baseline-climate scenario. Figure 1b does the same for the high-climate scenario.

Figure 1a. Dynamic costs of climate change in the baseline-climate scenario.





Figure 1b. Dynamic costs of climate change in the high-climate scenario.



The cost of BAU climate change in figure 1 is the sum of climate-change impacts and adaptation costs, but in practice the latter are very small in comparison.⁹ Thus figure 1 traces the global damage costs of climate change simulated by PAGE2002. Two key features become apparent. The first is that the greater portion of undiscounted damages accrue in the second half of the modelling horizon, between 2100 and 2200. In the baseline-climate scenario, mean damages are 2.6% of GWP in 2100, but 11.3% in 2200. In the high-climate scenario, they are 2.9% in 2100 and 13.8% in 2200. This is due to the convexity assumed in the damage functions, as well as the increased risk of catastrophic impacts that emerges at higher temperatures, which tend not to be realised until after 2100. The second is uncertainty around the mean estimates. The 5-95% confidence interval becomes increasingly wide as time passes and by 2200 it spans a range from 2.9% to 29.4% of GWP in the baseline-climate scenario and from 2.9% to 35.2% of GWP in the high-climate scenario. Impacts are in all cases higher in the high-climate scenario, because the higher regional mean temperatures attained trigger greater damage in economic and non-economic sectors, as well as increasing the risk of catastrophic impacts.

We now present the discounted present value of these total damage costs in terms of BGE losses (table 2a, b). Table 2a presents estimates for different pure rates of time preference, holding the elasticity of marginal utility of consumption constant at our preferred value of 1.0. The range of values adopted for

⁹ For example, climate-change impacts are over 400-fold higher than adaptation costs in the baseline-climate scenario, summed over the whole modelling period

the pure rate of time preference runs from that used in Stern (2006), a ‘prescriptive’ approach, to that following from Pearce and Ulph (1999), a ‘descriptive’ approach. Table 2b presents estimates for different elasticities of marginal utility of consumption, this time given our preferred pure rate of time preference of 0.1% per annum. The range of values adopted spans the brief literature review in section two, from a reasonable lower bound of 1.0 to a reasonable upper bound of 2.0.

Table 1a. Estimates of the discounted total cost of BAU climate change, in terms of BGE losses, for different pure rates of time preference

Pure rate of time preference	BGE losses (% loss in current consumption per capita), mean and 5%-95%	
	Baseline climate	High climate
0.1	11.1 (2.4-27.7)	14.7 (2.7-33.0)
0.5	8.1 (1.7-20.4)	10.6 (2.0-24.4)
1.0	5.2 (1.2-13.2)	6.7 (1.3-16.0)
1.5	3.3 (0.7-8.5)	4.2 (0.8-10.1)

Table 1b. Estimates of the discounted total cost of BAU climate change, in terms of BGE losses, for different elasticities of marginal utility of consumption.

Elasticity of marginal utility of consumption	BGE losses (% loss in current consumption per capita), mean and 5%-95%	
	Baseline climate	High climate
1.0	11.1 (2.4-27.7)	14.7 (2.7-33.0)
1.25	8.7 (2.2-21.7)	12.1 (2.7-26.0)
1.5	6.5 (1.7-16.5)	10.2 (2.0-20.0)
2.0	3.6 (1.1-9.4)	9.2 (1.0-11.4)

Table 1a illustrates the sensitivity of our estimates to an increase in the pure rate of time preference. A low pure rate of time preference, δ , yields a high present value of the total cost of BAU climate change. In the baseline-climate scenario, the mean cost equates to a permanent 11.1% fall in global mean per-capita consumption today (more precisely in 2001). The 5-95% confidence interval is 2.4-27.7%. In the high-climate scenario, the mean cost equates to a permanent 14.7% loss, with a 5-95% confidence interval of 2.7-33.0%. These are equivalent to the estimates reported in Stern (2006, chapter 6), with small differences generated by separate sets of Monte Carlo simulation draws. Increasing the pure rate of time preference to 0.5% p.a., 1.0% p.a. and 1.5% p.a. has a strong effect on estimated BGE losses. In the baseline-climate scenario, the mean cost falls steadily to a permanent loss of just 3.3% of global mean consumption per capita for $\delta=1.5\%$. The 5-95% confidence interval is 0.7-8.5%. Similarly, at $\delta=1.5\%$, the present value of the total cost of climate change in the high-climate scenario equates to just a 4.2% permanent loss, with a 5-95% confidence interval of 0.8-10.1%.

Table 2b demonstrates that increasing the elasticity of marginal utility of consumption from 1 to 2 also drives total cost estimates downwards. Where $\eta=2$, the mean BGE loss is 3.6% in the baseline-climate scenario and 9.2% in the high-climate scenario. This result is intuitive, since increases in η increase the social discount rate, but could not have been predicted *a priori*, since a higher elasticity also increases aversion to risk and therefore increases the relative weight apportioned to model runs with severe costs

of climate change. This is evident in the result that BGE losses fall more slowly in the high-climate scenario than in the baseline-climate scenario, since the high-climate scenario produces more runs with severe climate impacts. Nevertheless, the overall pattern of sensitivity tells us that the convexity in the overall damages estimated by PAGE2002 is insufficient to render aversion to risk as important as consumption smoothing in the overall estimate of expected utility. Increasing the damage function exponent in PAGE2002 would presumably increase the relative importance of risk aversion, a topic for future research and illustrated in the postscript to Stern (2006).

Comparing BGE estimates with dynamic estimates of the undiscounted cost of climate change shows the importance to the BGE aggregation methodology of two factors: first, model simulation draws that predict very low consumption; second, impacts in the latter part of the modelling period, together with what is assumed to happen after 2200 and how they both depend on the δ and η . We can directly scale losses of GWP to losses in consumption, since for convenience we assume a constant, exogenous saving rate. Where $\delta=0.1\%$ p.a. and $\eta=1$, the mean loss in undiscounted consumption in 2200 in the baseline-climate scenario is 11.3%, while the BGE loss is 11.1%. This is superimposed in figure 1a, as is the mean BGE loss in the high-climate scenario (figure 1b). It is clear that where both δ and η are at the low end of their ranges, our summary measure of total cost places much of its weight on the upper tail of the distribution of damages and on damages in the far-off future, including hysteresis in impacts beyond 2200. The importance of the latter decreases strongly as δ and η are increased, illustrated in tables 2a and b. The importance of the former increases, all else equal, when η is increased, but this is more than offset overall by the boost that increasing η gives to the social discount rate.

The effect of taking into account uncertainty

We also use the BGE metric of total cost to illustrate the effect of taking uncertainty into account in modelling the impacts of climate change. Table 3 presents estimates from four modelling approaches. The first is to set all stochastic parameters in PAGE2002 to their mode values. This simulates a modelling strategy in which best guesses are made about the value of all parameters, ‘best’ in the sense of being most likely. The second is to set all stochastic parameters in PAGE2002 to their mean values. This also simulates a deterministic strategy, but on this occasion we take some heed of cumulative asymmetry in the density functions describing the parameters. The third utilises the standard stochastic specification of PAGE2002, but instead of carrying out full expected-utility analysis as in tables 2a and b, we calculate the BGE for expected consumption, the mean path of global consumption per capita net of climate-change impacts and adaptation costs. The fourth reproduces the analysis in tables 2a and b, solving for the BGE equivalent to the expected utility of consumption. The analysis is, for now,

confined to the baseline-climate scenario, but reports results for different pure rates of time preference and different elasticities of marginal utility of consumption.

Table 3. Estimates of the total discounted cost of BAU climate change in the baseline scenario, under different modelling strategies and with different assumptions about discounting.

		<i>Elasticity of marginal utility of consumption</i>		
			<i>1.0</i>	<i>2.0</i>
		<i>Modelling strategy</i>	<i>BGE loss</i>	<i>BGE loss</i>
<i>Pure rate of time preference</i>	<i>0.1</i>	All modes	3.5	1.2
		All means	8.0	2.5
		Expected consumption	10.4	3.3
		Expected utility (mean loss)	11.1	3.6
	<i>1.5</i>	All modes	1.1	0.6
		All means	2.3	0.8
		Expected consumption	3.2	1.1
		Expected utility (mean loss)	3.3	1.1

From table 3 it is clear that a modelling strategy based on taking mode values for all parameters – equivalent to a best-guess strategy – leads to a significant underestimate of the total discounted cost of climate change, relative to an expected-utility approach. Where $\delta=0.1\%$ p.a. and $\eta=1$, this strategy yields an estimate of the BGE loss that is 3.5% of current global mean consumption per capita, relative to a mean of 11.1% when expected-utility analysis is applied. Moving to an ‘all means’ strategy reduces the shortfall significantly, although where $\delta=0.1\%$ p.a. and $\eta=1$, for example, it is still 3.1 percentage points below at 8%, a shortfall of 28% as a share of the expected-utility estimate. Calculating the utility of expected consumption brings the total discounted cost estimate much closer – to 10.4% where $\delta=0.1\%$ p.a. and $\eta=1$.

Increasing the pure rate of time preference and the elasticity of marginal utility of consumption reduces total cost estimates on all modelling strategies, because, as discussed above, both work to increase the social discount rate. The relative differences between the modelling strategies remain similar, except where both δ and η are increased together. Here, there is a reduction in the underestimate produced by adopting the ‘all modes’ strategy. Where $\delta=1.5\%$ p.a. and $\eta=2$, the BGE loss is 0.6%, which is 0.5 percentage points below the mean estimate from expected-utility analysis, or 45% relative to calculating expected utility. This compares with a relative underestimate of 68% (as a percentage of the expected-utility estimate) where $\delta=0.1\%$ p.a. and $\eta=1$. The explanation for this convergence must lie in the large increase in average social discount rates. Because impacts between 2100 and 2200 are valued much less, the effect of systematically incorporating uncertainty is also muted.

We might have expected that, where $\delta=0.1\%$ p.a. and $\eta=2$, there would be an increase in the relative difference between estimates that follow from taking the utility of expected consumption, and from taking the expected utility of consumption. Increasing the elasticity of marginal utility of consumption,

all else equal, increases the relative valuation on low-consumption simulation draws in expected-utility analysis. However, this is hardly observed in the data. Taking the utility of expected consumption underestimates the total discounted cost of climate change by 0.3 percentage points, or about 8% relative to calculating expected utility, while in the case where $\eta=1$ the difference is 0.7 percentage points or about 6%. In order to explore this result slightly further, we can look at the difference between these two modelling strategies in the high-climate scenario, for different values of η (table 4).

Table 4. Comparing expected consumption and expected utility approaches for the high-climate scenario.

	<i>Elasticity of marginal utility of consumption</i>	
	<i>1.0</i>	<i>2.0</i>
Expected consumption	12.7	4.0
Expected utility (mean loss)	14.4	9.2

The results for the high-climate scenario support our intuition much better. In this case, the underestimate produced by taking the utility of expected consumption is higher, and it is much higher where $\eta=2$. The difference between strategies here is 5.2 percentage points or 57% relative to the expected-utility approach. Figure 1 showed that the high-climate scenario produced a wider confidence interval, in particular with more low-consumption simulation draws. As a result, the high relative valuation of these draws made during expected-utility analysis drives the estimated BGE loss much higher than that calculated from an ‘expected consumption’ approach.

Discount factors

We now move on to consider the implications for discounting of uncertainty in estimating the impacts of climate change. We begin by demonstrating the differences in discount factors that are generated when climate-change impacts depress consumption. Figures 2a and b plot the difference between endogenous discount factors through time in the baseline- and high-climate scenarios, and the discount factors that are calculated by assuming consumption growth is unaffected by climate change. For convenience we label the latter discount factors ‘without climate change’. The latter is a typical modelling approach, although in many other cases the discount rate applied bears no actual relation to the rate of consumption growth assumed for the purposes of emissions growth, damage calculations and so on. We show differences with respect to the endogenously calculated schedule of certainty-equivalent discount factors, as well as the 5-95% range. The vertical axis is normalised to express differences as a percentage of the ‘without climate change’ discount factors.

Figure 2a. Relative differences between discount factors (as a % of ‘without climate change’ discount factors) for the baseline-climate scenario.





Figure 2b. Relative differences between discount factors (as a % of ‘without climate change’ discount factors) for the high-climate scenario.



Figures 2a and b demonstrate that significant relative differences emerge between the sets of discount factors. Looking first at the baseline-climate scenario in figure 2a, the certainty-equivalent discount factor is higher than the discount factor ‘without climate change’ throughout the modelling period, and it increases more than proportionately, again reflecting the assumption of overall convexity in climate

damages as temperatures increase. By 2200, it is around 14% higher, while the 5-95% range in 2200 is about 2-42%. To put this in context, the absolute certainty-equivalent discount factor is 0.07 in 2200, while the ‘without climate change’ discount factor is 0.06. As we would expect, the percentage difference is higher in the high-climate scenario, and the convexity in the relationship between the certainty-equivalent discount factors and the ‘without climate change’ discount factors is stronger. By 2200, the percentage difference is 50%, with a 5-95% range of 3-54%. The absolute certainty equivalent discount factor is 0.09 in 2200, again compared to 0.06 ‘without climate change’. The kink in the set of certainty-equivalent discount factors is due to linear interpolation of the PAGE2002 model outputs, which are available only for a limited number of years.

Next we examine what effect different discounting strategies have on estimates of the cost of climate change. We compare estimates of the marginal damage cost (MDC) of carbon with and without endogenous discounting (figure 3a-d). In the former case, we use the method derived from expected-utility analysis to calculate the MDC, outlined in section 2. In the latter case, we take the simplified route of summing the undiscounted marginal costs of climate change through time and discounting them based on a single rate, derived from consumption growth without climate-change impacts.¹⁰ In both cases, costs are measured for the modelling period, 2001-2200. We can also go some way towards illustrating the non-marginal nature of the climate-change problem by calculating the MDC of carbon along BAU trajectories, compared with stabilisation trajectories that control global warming. Results are presented for the baseline- and high-climate scenarios, and for stabilisation at 650ppm and 550ppm carbon dioxide equivalents (i.e. we stabilise the atmospheric concentration of all GHGs in the model, not just carbon dioxide). Since the stabilisation scenarios assume baseline climate change, they should not be directly compared with the high-climate scenario. That is, we cannot estimate the benefits of avoided climate change by comparing the MDC on the high-climate scenario with the MDC on a stabilisation scenario, because the stabilisation scenario assumes baseline climate change. To do so, we would have to model stabilisation assuming high climate change.

Figure 3a presents MDC estimates where $\delta=0.1\%$ p.a. and $\eta=1$. As we would expect, the MDC calculated through discounting endogenously is higher than the MDC calculated through discounting exogenously – i.e. based on assuming consumption is unaffected by climate change – for all four warming scenarios. What is particularly noteworthy, however, is that the error incurred by discounting exogenously becomes greater, the higher is the warming, in particular comparing the BAU scenarios with the stabilisation scenarios. In estimating the MDC on the high-climate scenario, discounting

¹⁰ We should, however, note that in the basic specification of PAGE2002, which calculates the MDC by discounting exogenously, the discount rate is based on regional consumption growth (thus giving rise to one exogenous discount rate per region), while our approach aggregates regional consumption to global consumption before discounting.

endogenously it is \$462.83/tC , compared with \$329.81/tC discounting exogenously. The error is about 30%, relative to the MDC discounting endogenously. On the baseline-climate scenario, the error is again about 30%. However, on the 650ppm and 550ppm stabilisation scenarios, the error is approximately 11%. Thus discounting exogenously masks the differences in MDCs across warming scenarios. The explanation for the overall pattern lies in the difference between the endogenously calculated discount factors and the set of discount factors ‘without climate change’. The higher the damages suffered in a particular warming scenario, the more these two sets will diverge. The former will rise further above the set of discount factors ‘without climate change’, and the error inculcated by discounting exogenously will increase. The MDC estimates diverge even on the stabilisation scenarios, because there are climate-change damage costs associated with these scenarios.

The overall pattern is repeated across figures 3b-d, which apply different combinations of the pure rate of time preference and the elasticity of marginal utility of consumption. There are two interesting results. Figure 3b presents results where $\delta=0.1\%$ p.a. and $\eta=2$. As risk aversion increases, the endogenously calculated MDC on the high-climate scenario rises significantly relative to its counterpart discounted exogenously, and relative to both sets of MDCs on the other warming scenarios. Discounting endogenously, it is \$155.93/tC, lower than where $\eta=1$ due to the lower consumption discount factors, but discounting exogenously it is just \$52.07/tC, an error of 67% relative to the former. By contrast, figure 3c shows that increasing the pure rate of time preference compresses the error caused by discounting exogenously. On the high-climate scenario, the error is just 21%, and on the 550ppm stabilisation scenario it is 5%. This is because higher utility discounting reduces the relative importance of low-consumption growth later in the modelling period.

Figure 3a. Difference between marginal damage cost estimates for endogenous and exogenous discounting, $\delta=0.1\%$ p.a., $\eta=1$.




Figure 3b. Difference between marginal damage cost estimates for endogenous and exogenous discounting, $\delta=0.1\%$ p.a., $\eta=2$.




Figure 3c. Difference between marginal damage cost estimates for endogenous and exogenous discounting, $\delta=1.5\%$ p.a., $\eta=1$.





Figure 3d. Difference between marginal damage cost estimates for endogenous and exogenous discounting, $\delta=1.5\%$ p.a., $\eta=2$.



At this stage, it is important to comment on the differences between MDCs on BAU scenarios and MDCs on stabilisation scenarios. That there could be a difference is intuitive, because we specify a convex relationship between damages and temperature, so that if a marginal tonne of carbon is added to a higher-temperature scenario – for example comparing the BAU scenario with stabilisation – then its MDC should accordingly be higher. On the other hand, Hope (2005), also using the PAGE2002 model, estimated that the MDC of one tonne of carbon emitted today is in fact insensitive to the path of emissions on which it is added. This analysis also included a comparison of the IPCC A2 scenario, as BAU, with 550ppm stabilisation. This apparently counter-intuitive result was shown to be due to the interplay between the logarithmic relationship linking radiative forcing and GHG concentrations, the convex relationship between damage and temperature, and discounting.

The explanation for this discrepancy could lie in a number of factors. First, we stabilise total GHG concentrations, rather than simply the concentration of carbon dioxide. As a result of emissions of other GHGs, total radiative forcing in Hope's (2005) 550ppm stabilisation scenario is significantly higher than 550ppm carbon dioxide equivalent, whereas our scenario, which controls other GHGs as well, stabilises mean forcing at around that level. In turn, our stabilisation scenario produces much less warming than Hope's earlier scenario, for example 2.8°C warming by 2200 on average, compared with 5.1°C warming on average in the earlier study. Second, one element of the 'success' of our stabilisation scenario, so-to-speak, is the assumption that the emission of sulphate aerosols, which have a cooling effect, continues to follow a BAU trend, while emissions of other GHGs are reduced strongly. Future research should examine the effect of reducing sulphate emissions, as it would be most likely to close the gap between MDCs on a BAU scenario and MDCs on a stabilisation scenario, perhaps strongly. Third, we of course use different methods of calculating the MDC, here using expected-utility analysis. Figure 3 illustrates that discounting exogenously in calculating the MDC will suppress some of the differences between paths. Fourth, we specify δ and η such that our average social discount rate is in almost all cases lower than in Hope (2005), and higher discounting of course reduces the value of convex damages in the more distant future.

Finally, we compare the probability distribution of discount factors that is calculated along a BAU warming scenario with that calculated along a global policy scenario, in which the choice is made to stabilise atmospheric GHG concentrations by making an up-front investment in emission abatement. Figures 4a and b compare certainty-equivalent discount factors on a BAU consumption scenario (baseline climate) with a 550ppm-stabilisation consumption scenario. In this case, the latter is net of the cost of reducing GHG emissions. Figure 4a presents the difference where $\delta=0.1\%$ p.a. and $\eta=1$, and given costs of mitigation that are around 1% of GWP throughout the modelling period (estimates by

Dennis Anderson for Stern, 2006). As we would expect, consumption on the policy scenario, as reflected in the path of certainty-equivalent discount factors, falls relative to the BAU scenario initially, as consumption is diverted into investment in GHG emission abatement. However, this pays back handsomely after around 2080, when climate-change impacts bring the BAU consumption scenario below the stabilisation scenario. By 2200, the difference in the certainty-equivalent discount factors is substantial: 10.9% relative to the BAU path. In absolute terms, the certainty-equivalent discount factor on the BAU path is 0.061 in 2200, while it is 0.069 on the 550ppm stabilisation path. It is clear that the same consumption scenario will give rise to a significant difference in present value cost estimates (total or marginal), depending on which set of discount factors is used, which illustrates neatly the non-marginal nature of the relationship. Of course, it would be more consistent to treat the costs of GHG emission reductions as stochastic, in the same way as the benefits are. This is another topic for further work.

Identical relative differences are returned if the pure rate of time preference is increased. However, the difference works on a much lower certainty-equivalent discount factor in absolute terms. Where $\delta=1.5\%$ p.a., the certainty-equivalent discount factor on the BAU scenario is 0.0042 in 2200, while it is 0.0038 on the 550ppm stabilisation scenario. Further research will investigate the effect of increasing the elasticity of marginal utility of consumption. Although we know that the rate of fall of both sets of certainty-equivalent discount factors will increase, the relative difference may also increase, due to the relative valuation of low-consumption paths on the BAU scenario, which are avoided by stabilisation.

In figure 4b, we investigate the impact of increasing costs of mitigation. In this case, costs increase from 0-4% of GWP between around 2001 and 2080, thereafter holding at 4% (data from Dennis Anderson). Figure 4b shows that the certainty-equivalent discount factors for 550ppm stabilisation fall further below the BAU scenario to begin with, reaching a maximum of around 3.5% below the BAU by around 2075. However, consumption on the policy scenario becomes greater than that on the BAU scenario after around 2125, and increases until the certainty-equivalent discount factor is 8.1% higher relative to the BAU in 2200. In absolute terms, the certainty-equivalent discount factor on the BAU path is 0.063 in 2200, while it is 0.069 on the 550ppm stabilisation path. If we were to calculate the average social discount rate from 2001 to 2200 in this case, the difference between the two scenarios would be smaller than it would be for figure 4a, because the foregone consumption early in the modelling period due to more costly investment in GHG emission abatement balances out the gain in consumption due to avoided climate impacts later on.

Figure 4a. Relative differences between certainty-equivalent discount factors on a BAU scenario (baseline climate) and a 550ppm stabilisation scenario (as a % of BAU), with mitigation costs around 1% of GWP, $\delta=0.1\%$ p.a. and $\eta=1$.





Figure 4b. Relative differences between certainty-equivalent discount factors on a BAU scenario (baseline climate) and a 550ppm stabilisation scenario (as a % of BAU), with mitigation costs around 4% of GWP, $\delta=0.1\%$ p.a. and $\eta=1$.



5. Discussion

The science of climate change dictates the nature of the challenge for economics. In so doing, it raises the possibility, however small, of large, potentially irreversible costs on the global economy and on welfare, brought about when warming generates high global mean temperatures relative to pre-industrial levels. For theorists and practitioners of applied CBA, it dictates that we develop techniques of aggregation across time, space and possible states of the world that are able to systematically, consistently and fully take into account wide confidence margins and large differences in consumption.

In this paper, we have sought to clearly state an expected-utility approach to CBA of climate-change policies, and to systematically evaluate the consequences of adopting such an approach, in comparison with the simplified techniques of CBA developed for project appraisal. We are aware that some previous papers have considered elements of this issue (e.g. Tol, 1999, 2003), but we focus more strongly on a comparative treatment of cost-benefit approaches. There remains much research to do here, for example in incorporating aggregation across space, in separating assumed preferences for consumption smoothing from risk aversion, and in estimating stochastic costs of mitigation. But we trust that we have demonstrated sufficiently the promise and importance of our general approach.

Our empirical results demonstrate the consequences of failing to adequately take uncertainty into account. Table 3 shows that the specification of a deterministic IAM can lead to a large underestimate of the total cost of climate change. This is especially true if one specifies an ‘all modes’ model, in which ‘best guesses’ are made about all parameter values, taking as a best guess the mode of every density function that describes uncertainty in the underlying literature. Even a slightly more sophisticated ‘expected consumption’ approach, which uses a probabilistic model to estimate average consumption, and then computes the utility of that consumption, could lead to a large underestimate, if the confidence interval describing climate-change impacts is especially wide and skewed to the side of heavy losses, as is the case in our high-climate scenario.

Figures 2, 3 and 4 go on to illustrate that the wide range of potential climate-change impacts strongly affects the magnitude and rate of change of the social discount factor. What is more, failing to discount endogenously, based on growth in consumption that is actually realised in a world with climate change, leads to systematic and in some cases substantial underestimation of the marginal damage cost of carbon. The relative underestimation increases, the higher is the elasticity of marginal utility of consumption, because model simulation draws in which consumption is especially low are valued more highly, all else equal. Increasing the pure rate of time preference mutes this underestimation, because larger impacts further into the future are valued less. Figure 4 goes on to plot certainty-equivalent discount factors on a BAU path, relative to a global policy path, in which substantial investment in

emission abatement is made early on, to deliver higher consumption later. Again, the discount factors are substantially different, and we may infer that failure to use the appropriate set will generate a substantial error.

Table 2 shows that both the pure rate of time preference and the elasticity of marginal utility of consumption have a strong bearing on estimates of the total cost of climate change. Both work to increase the social discount rate, so our estimate of the total cost of BAU climate change in the future falls quite markedly. Where both are at the low end of the range we have set, the mean total cost of climate change, in terms of a permanent loss in global mean consumption per capita today, is strikingly high, being 11.1% in the baseline-climate scenario and 14.7% in the high-climate scenario. In both scenarios, this lies close to the maximum undiscounted mean cost of climate change in 2200, which tells us that what we assume about impacts later in the modelling period, and beyond, matters much to our ‘non-marginal’ method. But, in commenting on what our analysis suggests about the seriousness of the climate-change externality, I note that for social rates of discount that are supported by a (disputed) examination of ethics and empirics, the total cost is conclusively high. What happens after 2200 is important, but in closing off the modelling period at 2200, we are arguably underestimating the capacity of climate change to continue depressing consumption further and further below its counterfactual level. Similarly, I see hysteresis as the most reasonable assumption at the end of modelling period, for want of better evidence. It is difficult to see any persuasive reason why successively lower output trajectories, with lower consumption and lower investment, would have the same steady-state in the very long-run. By 2200, for instance, most of the total damage estimated by PAGE2002 is working through non-economic impacts on human health and ecosystems, as well as the unfolding risk of a global climate discontinuity (Warren *et al.*, 2006). Neither of these types of impact can plausibly be reversed. Learning to live with them is one thing, and is indeed what we implicitly assume: using them as a spur to growth in consumption that is faster than it would otherwise have been is quite another.

That the MDC of carbon may vary across emission scenarios, in particular between BAU scenarios and stabilisation scenarios, is an important preliminary result. I note the discord between results here and in Hope (2005), but if future research is able to confirm this difference, using expected-utility analysis to liberate the differences between marginal costs, then there are tricky implications for policy, plan, programme and project appraisals that would use the MDC of carbon as a benchmark. In theory, it is simple enough to say that we should apply the MDC corresponding to the path that we are, in fact, on, worldwide. But this is difficult to determine, especially over time, as we presumably become more and more carbon constrained. Perhaps for now a pragmatic approach is to subject any appraisal to a range of MDCs, spanning, amongst other uncertainties, uncertainty over future global policy. The downside is

that ranges are notoriously difficult to use in project appraisal: what happens, for instance, if a project passes a cost-benefit test on a business-as-usual MDC but not on a stabilisation MDC? Clearly some sophisticated, but practicable, rules of thumb are required.

Finally, it leaves me to comment on how far we can currently push CBA as a tool for determining global climate-change policy. In section one I argued that there would always be an in-principle need for CBA, as one member of a range of decision-support methodologies. At the same time, I argued that it was not straightforward to dismiss CBA, even if some of the impacts of climate change work on non-substitutable natural capital. We are effectively modelling the loss of critical natural assets in those IAMs where there is a risk of catastrophic impacts. But everything in the end turns on our ability to define subjective probabilities over these 'risks'. If there remains unquantified uncertainty, then it is logical to conclude that CBA can only be applied to a subset of the potential impacts that we are now contemplating. This seems to be the situation here. Looking for instance at warming in our high-climate scenario, it lies well within the extremes reported in Meinshausen's (2006) synthesis of eleven studies. And the method on which we rely to calibrate the risk of catastrophic impacts is currently crude indeed. Researching for their 2000 book, Nordhaus and Boyer polled a number of experts on the probability that a very large loss of 25% of global GDP, roughly equivalent to the effect of the Great Depression, would result from increases in global mean temperature of 3°C by 2090, 6°C by 2175 and 6°C by 2090. Taking account of estimated differences in regional vulnerability to catastrophic climate change, they used survey data to estimate people's willingness to pay to avoid the resulting risk. PAGE2002 is broadly calibrated on this evidence.

Tol (2003) considers the further complication encountered if uncertainty (variance to be precise) is very large, perhaps infinite. He notes that, for expected CBA to be applicable, both the costs and benefits of emission reductions must exhibit finite variance. His IAM, FUND, is specified to exhibit finite variance in its parameters (like PAGE2002), but the variance in the expected MDC can become infinite – and the expected MDC itself undefined – if consumption growth becomes persistently negative due to climate change. Discounting endogenously, the discount factors on one of his Monte Carlo simulation draws become infinite. If one discounted exogenously, this outcome is unlikely to obtain. This is not an argument against expected-utility analysis, for all the reasons we have rehearsed in this paper, but rather raises questions about the nature of uncertainty and the role of CBA. Tol does point out, however, that infinite variance exists both with and without policy to reduce GHG emission abatement. Thus the pertinent question is whether uncertainty in the difference between the scenarios is finite? This is difficult to prove but more likely. In the mean time, we constrain the variance in the expected MDC estimated in this paper, to produce tractable results.

So we are not yet in a position to place great faith in the outcomes of CBA of climate-change policy, even if we all agree to be welfarists. Ultimately, the pragmatic approach that is evident for example in the IPCC's 'Reasons for Concern' remains reasonable overall.

References

- Arrow, K. J. (1995). Inter-generational equity and the rate of discount in long-term social investment. IEA World Congress, December 1995.
- Asheim, G. B. and W. Buchholz (2003). "The malleability of undiscounted utilitarianism as a criterion of intergenerational justice." Economica **70**: 405-422.
- Azar, C. (1998). "Are optimal emissions really optimal? Four critical issues for economists in the greenhouse." Environment and Resource Economics **13**: 249-268.
- Azar, C. and T. Sterner (1996). "Discounting and distributional considerations in the context of global warming." Ecological Economics **19**(2): 169-184.
- Barry, B. (1991). Liberty and Justice: Essays in Political Theory 2. Oxford, Clarendon Press.
- Broome, J. (1992). Counting the Cost of Global Warming. Cambridge, White Horse Press.
- Cline, W. R. (1992). The Economics of Global Warming. Washington, Institute for International Economics.
- Cowell, F. A. and K. Gardiner (1999). Welfare Weights. Economics Research Paper 20, STICERD, London School of Economics and Political Science.
- Daly, H. E. (1992). Steady-State Economics. London, Earthscan.
- Daly, H. E. (1996). Beyond Growth. Boston, Beacon Press.
- Daly, H. E. and R. Costanza (1992). "Natural capital and sustainable development." Conservation Biology **6**(1): 37-46.
- Epstein, L. G. and S. E. Zin (1989). "Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework." Econometrica **57**(4): 937-969.
- Fankhauser, S. (1995). Valuing Climate Change: the Economics of the Greenhouse. London, Earthscan.
- Friedlingstein, P., P. Cox, et al. (2006). "Climate-carbon cycle feedback analysis: results from C4MIP model intercomparison." Journal of Climate **19**: 3337-3353.
- Gedney, N., P. Cox, et al. (2004). "Climate feedback from wetland methane emissions." Geophysical Research Letters **31**(20): L20503.
- Gerlagh, R. and B. C. C. van der Zwaan (2002). "Long-term substitutability between environmental and man-made goods." Journal of Environmental Economics and Management **44**: 329-345.
- Gollier, C. (2001). The Economics of Risk and Time. Cambridge, Mass, MIT Press.
- Gollier, C. (2002a). "Time horizon and the discount rate." Journal of Economic Theory **107**(2): 463-473.
- Gollier, C. (2002b). "Discounting an uncertain future." Journal of Public Economics **85**: 149-166.
- Groom, B., C. J. Hepburn, et al. (2005). "Discounting the future: the long and the short of it." Environment and Resource Economics **31**(1): 445-493.
- Hepburn, C. J. (2006). Discounting Climate Change Damages: Working Note for the Stern Review.

- Hope, C. (2005). The economics of climate change. House of Lords Select Committee on Economic Affairs, Volume II: evidence. London, The Stationery Office.
- Hope, C. (2006). "The marginal impact of CO₂ from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern." Integrated Assessment **6**(1): 19-56.
- Hubin, D. C. (1994). "The moral justification of benefit/cost analysis." Economics and Philosophy **10**: 169-194.
- IPCC (2001). Summary for Policymakers: A Report of Working Group I of the Intergovernmental Panel on Climate Change. Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change. J. T. Houghton, Y. Ding, D. J. Griggset al. Cambridge, Cambridge University Press: 1-20.
- Knight, F. H. (1921). Risk, Uncertainty and Profit. Boston, Houghton Mifflin.
- Kreps, D. M. and E. L. Porteus (1978). "Temporal resolution of uncertainty and dynamic choice theory." Econometrica **46**(1): 185-200.
- Maddison, D. J. (1995). "A Cost-Benefit Analysis of Slowing Climate Change." Energy Policy **23**(4/5): 337-346.
- Manne, A. S., R. Mendelsohn, et al. (1995). "MERGE - a model for evaluating regional and global effects of GHG reduction policies." Energy Policy **23**(1): 17-34.
- Mastrandrea, M. D. and S. H. Schneider (2004). "Probabilistic integrated assessment of "dangerous" climate change." Science **304**: 571-575.
- McCarthy, J. J., O. F. Canziani, et al., Eds. (2001). Climate Change 2001: Impacts, Adaptation, and Vulnerability: Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK, Cambridge University Press.
- Meinshausen, M. (2006). What does a 2degC target mean for greenhouse gas concentrations? A brief analysis based on multi-gas emission pathways and several climate sensitivity uncertainty estimates. Avoiding Dangerous Climate Change. H.-J. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley and G. Yohe. Cambridge, UK, Cambridge University Press.
- Mirrlees, J. A. and N. H. Stern (1972). "Fairly good plans." Journal of Economic Theory **4**(2): 268-288.
- Nakicenovic, N. and R. Swart, Eds. (2000). Special Report on Emissions Scenarios: a Special Report of Working Group III of the Intergovernmental Panel on Climate Change. Cambridge, Cambridge University Press.
- Neumayer, E. (1999). "Global warming: discounting is not the issue, but substitutability is." Energy Policy **27**: 33-43.
- Newell, R. and W. Pizer (2003). "Discounting the benefits of climate change mitigation: how much do uncertain rates increase valuations." Journal of Environmental Economics and Management **46**(1): 53-71.
- Nordhaus, W. D. (1991). "To slow or not to slow: the economics of the greenhouse effect." Economic Journal **101**: 920-937.
- Nordhaus, W. D. (1994). Managing the Global Commons: the Economics of Climate Change. Cambridge, Mass, MIT Press.
- Nordhaus, W. D. (2006). A Review of the Stern Review on the Economics of Global Warming.
- Nordhaus, W. D. and J. Boyer (2000). Warming the World: Economic Models of Global Warming. Cambridge, Mass., MIT Press.
- OXERA (2002). A Social Time Preference Rate for Use in Long-Term Discounting. London, The Office of the Deputy Prime Minister, Department for Transport and Department of the Environment, Food and Rural Affairs.
- Pearce, D. W., B. Groom, et al. (2003). "Valuing the future: recent advances in social discounting." World Economics **4**(2): 121-141.

- Pearce, D. W. and D. Ulph (1999). A social discount rate for the United Kingdom. Environmental Economics: Essays in Ecological Economics and Sustainable Development. D. W. Pearce. Cheltenham, Edward Elgar: 268–285.
- Peck, S. C. and T. J. Teisberg (1992). "CETA: a model for carbon emissions trajectory assessment." Energy Journal **13**(1): 55-77.
- Plambeck, E. L. and C. Hope (1996). "PAGE95: an updated valuation of the impacts of global warming." Energy Policy **24**(9): 783-794.
- Plambeck, E. L., C. Hope, et al. (1997). "The PAGE95 model: integrating the science and economics of global warming." Energy Policy **19**: 77-101.
- Ramsey, F. P. (1928). "A mathematical theory of saving." Economic Journal **38**(543-559).
- Selden, L. R. (1979). "An OCE analysis of the effect of uncertainty on saving under risk preference independence." The Review of Economic Studies **46**(1): 73-82.
- Solow, R. M. (1974). "Intergenerational equity and exhaustible resources." Review of Economic Studies Symposium: 29-46.
- Spash, C. L. (1994). "Double CO2 and beyond: benefits, costs and compensation." Ecological Economics **10**(1): 27-36.
- Stern, N. (2006). *The Economics of Climate Change*. London, HM Treasury.
- Stern, N. H. (1977). The marginal valuation of income. Studies in Modern Economic Analysis. M. Artis and R. Nobay. Oxford, Blackwell.
- Tol, R. S. J. (1994). "The damage costs of climate change: a note on tangibles and intangibles, applied to DICE." Energy Policy **22**(5): 436-438.
- Tol, R. S. J. (1997). "On the optimal control of carbon dioxide emissions: an application of FUND." Environmental Modelling and Assessment **2**: 151-163.
- Tol, R. S. J. (1999). "Safe policies in an uncertain climate: an application of FUND." Global Environmental Change **9**: 221-232.
- Tol, R. S. J. (2005). "The marginal damage costs of carbon dioxide emissions: an assessment of the uncertainties." Energy Policy **33**(16): 2064-2074.
- von Neumann, J. and O. Morgenstern (1944). Theory of Games and Economic Behavior. Princeton, NJ, Princeton University Press.
- Warren, R., C. Hope, et al. (2006). Spotlighting impacts functions in integrated assessment: research report prepared for the Stern Review on the Economics of Climate Change. Working Paper 91, Tyndall Centre.
- Weitzman, M. (1998). "Why the far distant future should be discounted at its lowest possible rate." Journal of Environmental Economics and Management **36**: 201-208.
- Weitzman, M. (2004). Discounting a distant future whose technology is unknown. Cambridge, Mass, Harvard University, mimeo.
- Yohe, G. (2006). "Some thoughts on the damage estimates presented in the Stern Review: an editorial." Integrated Assessment **6**(3): 65-72.